Modeling, Data Assimilation and Advanced Computing

Earth System Research Laboratory

Mark Govett

Advanced Computing

NOAA HPC Resources and where Boulder fits

- System supports OAR science activities
  - Tolerates lower reliability
  - Better price-performance
- Proving ground for new technologies
  - Computer systems
  - Mass Storage
  - Networking
  - File systems

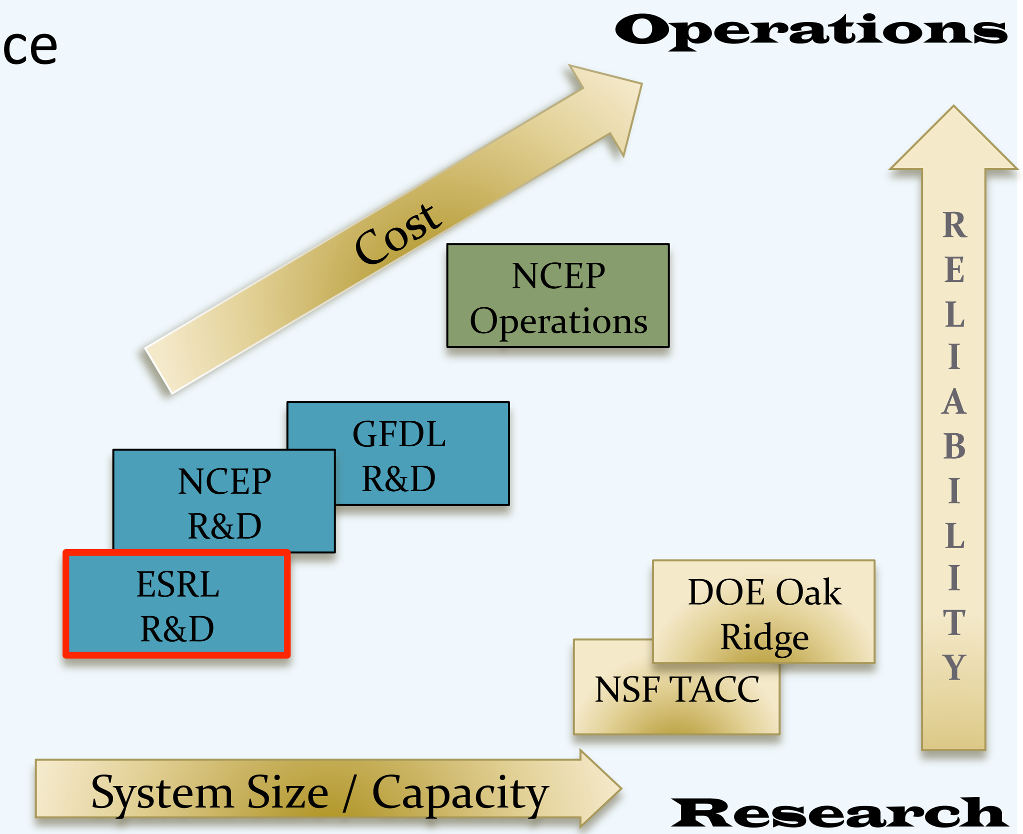

Computer Systems

Cost effective, cutting-edge, innovative

- 1992: Purchased NOAA's first Massively Parallel Processor (MPP) super-computer
  - Low-cost alternative to expensive vector machines
- 2000: First NOAA Linux cluster
  - Used Commercial Off-The-Shelf technology (COTS)
  - 8th fastest computer (2002: TOP500 list)
- 2008: Began exploring the use of Graphics Processing Units (GPUs)

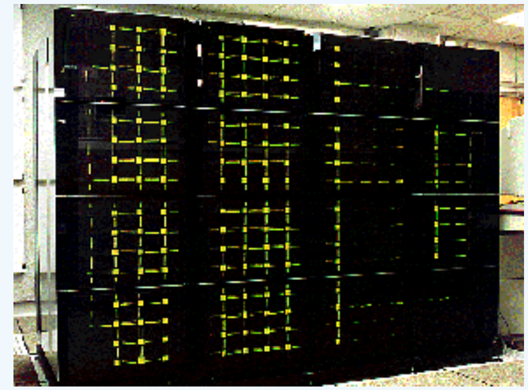

MPP System (1992)

Linux Cluster (2000)

PetaFlop Computing

- Required for global cloud resolving scales (3-4km)
  - Large CPU systems (~200 thousand cores) are unrealistic for operational weather forecasting
    - Power, cooling, reliability, cost
- CPU chip performance flat
- GPUs appear to be the future of HPC

GPU: 2008
933Gflops
150W

Fermi (2010)

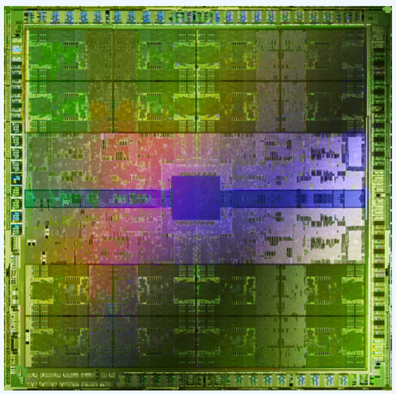

✧ 8x increase in double precision
✧ 2x increase in memory bandwidth
✧ Error correcting memory

Tesla (2008)

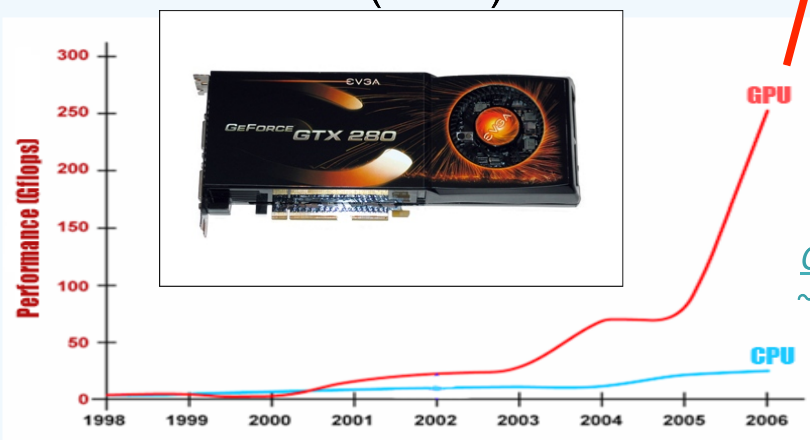

CPU:2008
~45 Gflops
160W

Software Development

from Vector to MPPs (1990s)

- Developed the Scalable Modeling System (SMS)
  - Directive-based code parallelization
  - Handle inter-process communications, I/O
    - Distributed memory paradigm
  - Demonstrated good performance and portability

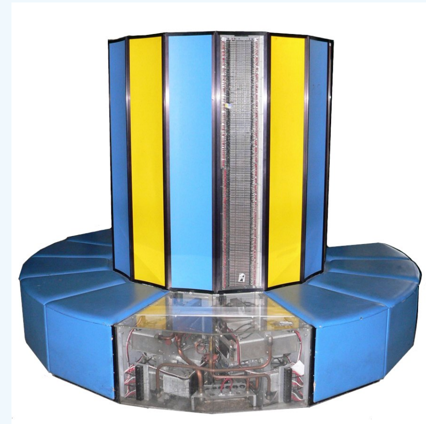

Cray XMP
Cost: $20-40M

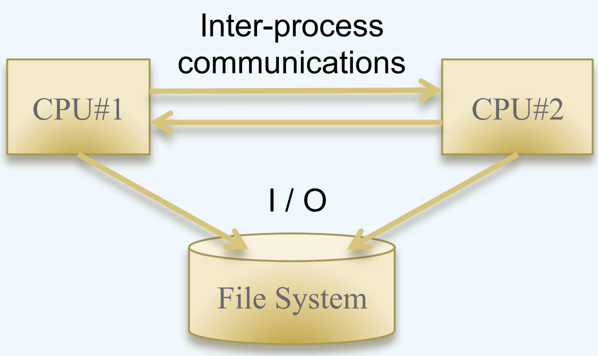

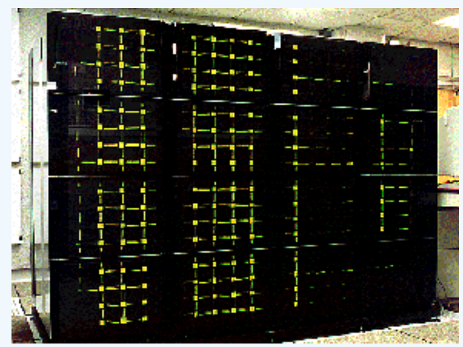

Intel Paragon (1992)
Cost: $1M

Outcomes

- Demonstrated MPPs were a cost-effective technology
- Led to NCEP moving away from expensive vector machines

Software Development

from MPPs to Linux Clusters (2000s)

- Worked with vendors to develop software infrastructure
  - Node management, file systems, batch systems

IBM Cluster (2008)
Vendor Solution

Outcomes

- Demonstrated Linux clusters could be assembled and managed at a huge cost savings
  - enables more science at Boulder facility
  - HRRR using 900 CPUs for hourly cycled runs
- Many large research systems are now Linux clusters

TACC Ranger (2008)
Linux Cluster
62,976 CPU cores
579 TFlops peak

Linux Cluster (2008)

COTS Solution
2-5 times cheaper

Software Development

from CPUs to GPUs (2010s)

- Purchased 16 node GPU system in 2008
- Developed Fortran to CUDA compiler
  - Commercial compilers now available
- Parallelized all of NIM dynamics
  - Runs 34 times faster on a GPU than CPU
    - Single GPU, Communications only for I/O

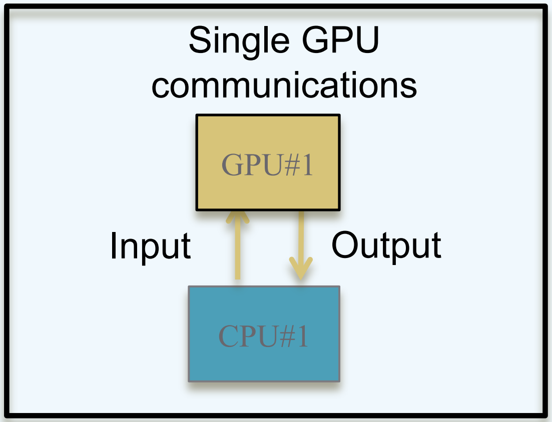

<u>Current Efforts</u>

- Run NIM dynamics on multiple GPU nodes
- Parallelize FIM, HYCOM, GFS physics for GPU

<u>Outcomes</u>

- Public release of Fortran-GPU compiler
  - Noted on NVIDIA site, technical papers
  - Downloads worldwide

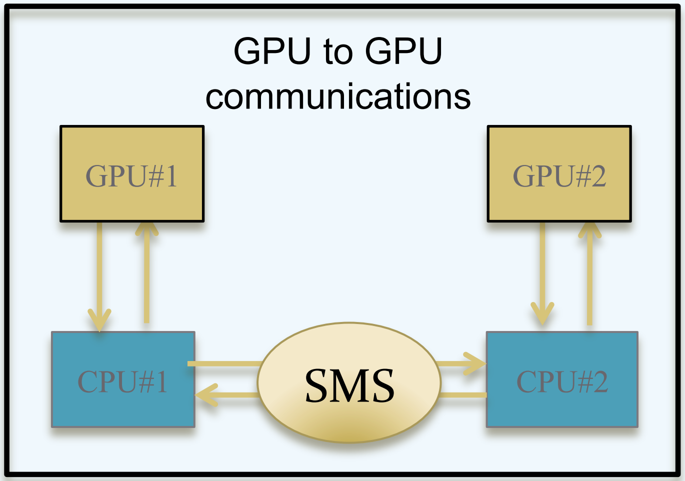

Models & Ensembles Software Investment

- Models are becoming increasingly complex
  - Increasing emphasis on sharing models, components
- Portability & performance important to NOAA
  - 2008 GSD port of FIM to TACC took 3 days (HFIP)
- Interoperability
  - Global Interoperability Program (GIP)

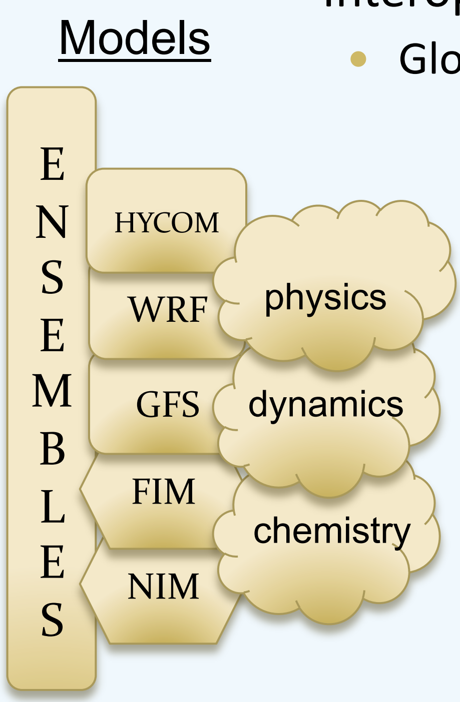

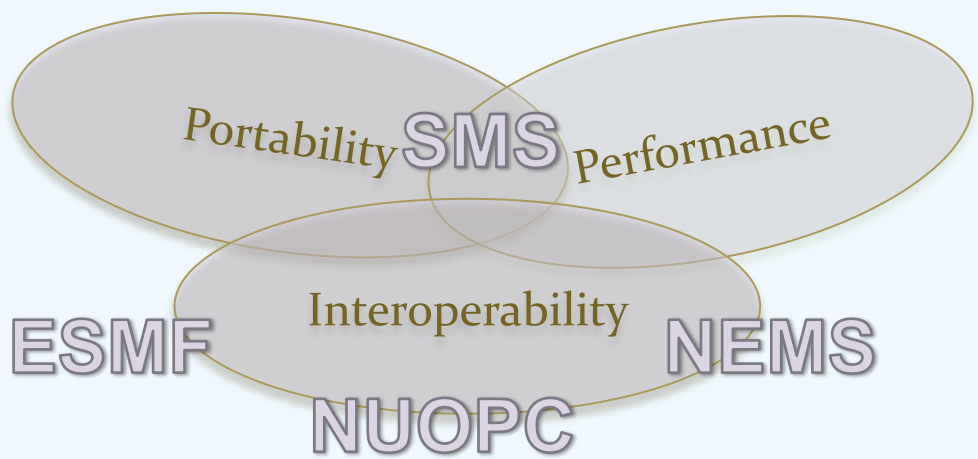

NOAA systems

- NCEP Operations
- GFDL R&D
- NCEP R&D
- Boulder R&D
- DOE Oak Ridge
- NSF TACC

Final Thoughts

- GSD plays a vital role for NOAA in Advanced Computing
  - Support OAR science with low-cost, high-performance, cutting-edge systems
  - Continue to provide leadership in HPC
    - Exploring new hardware technologies
    - Developing support tools
    - Enhancing model portability, performance and interoperability
  - These activities continue to be very beneficial for NOAA and the wider community